FINITENESS CONDITIONS ON TRANSLATION SURFACES

JOSHUA P. BOWMAN

Introduction

Throughout this note, let X denote a *translation surface*, i.e., a (connected) topological surface with a translation atlas. Then X is automatically endowed with a conformal structure and a flat metric, and so it is both a Riemann surface and a Riemannian manifold [HM]. An orientation-preserving homeomorphism $\phi : X \to X$ is called *affine* if it is affine in local charts. We use $\mathrm{Aff}^+(X)$ to denote the group of affine maps of X. Any element ϕ of $\mathrm{Aff}^+(X)$ has a well-defined global *derivative* $\mathrm{der}\,\phi \in \mathrm{GL}_2^+(\mathbb{R})$. The image $\Gamma(X)$ of the homomorphism $\mathrm{der} : \mathrm{Aff}^+(X) \to \mathrm{GL}_2^+(\mathbb{R})$ is called the *Veech group* of X [Ve, Vo, EG, GJ].

The existence of affine self-maps of a translation surface has applications in the study of mapping class groups, Teichmüller theory, algebraic geometry, and dynamical systems (for a small sampling of such applications, see, e.g., [Th, HS, Mc, Mö, LR, De]). They measure a kind of "symmetry" more general than that of isometries, which nonetheless has consequences for such systems as geodesic flow on the surface and geodesics in Teichmüller space. Veech first observed the importance of the group of derivatives of affine maps [Ve].

Let $\overline{X}$ denote the metric completion of X. The classical study of translation surfaces assumes that $\overline{X}$ is itself a compact surface and $\overline{X} \setminus X$ is finite. If these conditions are satisfied, we will say that X has *finite affine type*. Here we wish to consider four other "finiteness" conditions that may be placed on X:

(1) X has *finite analytic type* as a Riemann surface, meaning that it is obtained from a compact Riemann surface by making finitely many punctures.
(2) X has *finite area* as a Riemannian manifold, meaning that the integral of the induced area form over all of X is finite.
(3) X is *bounded* as a metric space, meaning that there exists a constant $M > 0$ such that $d_X(x, y) \leq M$ for every pair of points x and y in X.
(4) X is *totally bounded* as a metric space, meaning that for any fixed $\varepsilon > 0$, X can be covered by finitely many balls of radius ε (equivalently, $\overline{X}$ is compact).

We will prove two main results about these conditions, one negative and one positive.

Theorem 1. *Except for the trivial implication "totally bounded $\implies$ bounded", none of the conditions* (1)–(4) *on X implies any of the others. However, if X has finite analytic type, then the other three conditions are equivalent and imply that X has finite affine type.*

Theorem 2. *Suppose the ideal boundary of X is empty. If X has at least one periodic trajectory and is totally bounded or has finite area, then its Veech group is a discrete subgroup of $\mathrm{SL}_2(\mathbb{R})$. However, there exist bounded surfaces and surfaces of finite analytic type with non-discrete Veech groups.*

2010 *Mathematics Subject Classification.* Primary 30F30. Secondary 57M50, 32G15, 14K20.

Remark. It is likely that the condition of having a periodic trajectory follows from the assumptions of having empty ideal boundary and being totally bounded or of finite area, in which case it can be dropped in Theorem 2.

Translation surfaces of infinite analytic type appear, for example, in [EG, CGL, HLT, HHW, Va, Bo], and it is such examples that motivated the study presented here. We will prove Theorem 1 in §1 and Theorem 2 in §2.

1. Inequivalence of finiteness conditions

We begin with the trivial, and only, implication among the finiteness conditions (1)–(4).

Proposition 1.1. *"X is totally bounded" $\implies$ "X is bounded".*

Proof. This is a generality about metric spaces. Pick $\varepsilon > 0$, and cover X with N balls of radius ε. Then the distance between any two points is at most $2N\varepsilon$. □

The rest of the first part of Theorem 1 is proved through a series of examples. One general construction will be quite useful and flexible, so we describe it first and establish some notation.

Example 1 (An infinite "stack of boxes"). Let $H = \{h_n\}_{n=1}^{\infty}$ be a sequence of positive numbers, and let $W = \{w_n\}_{n=1}^{\infty}$ be a strictly decreasing sequence of positive numbers tending to zero. Then we construct a translation surface $X_{H,W}$ as follows (see Figure 1):

- For each $n \geq 1$, let R_n be a rectangle with horizontal side w_n and vertical side h_n.
- Place the sequence of rectangles in the plane $\mathbb{R}^2$, starting with R_1 having its lower left corner at the origin, and with R_{n+1} immediately above R_n so that its left edge is along the y-axis.
- Identify the right and left sides of each R_n with each other via horizontal translation, and identify the portion of the top of R_n not covered by R_{n+1} (of length $w_n - w_{n+1}$) with the portion of the bottom edge of R_1 directly below via vertical translation. (We omit the vertices.)

The genus of $X_{W,H}$ is infinite, as can be seen by considering the (pairwise non-homotopic) horizontal core curves of the R_n. The area of $X_{H,W}$ is

$$\mathrm{Area}(X_{H,W}) = \sum_{n=1}^{\infty} \mathrm{Area}(R_n) = \sum_{n=1}^{\infty} h_n w_n.$$

In particular, the area of $X_{H,W}$ is finite if H and W are sequences in ℓ^2, but this is not necessary. Let $\overline{X}_{H,W}$ denote the metric completion of this surface.

Lemma 1.1. *$\overline{X}_{H,W} \setminus X_{H,W}$ has only one point.*

Proof. The translation structure has been defined by taking a quotient of the union of the rectangles R_n except for their vertices. The vertices are all collapsed to a single point, as is evident in Figure 1: using the notation of that figure, observe that $A_0 \sim A_1 \sim B_1 \sim C_1 \sim D_1 \sim A_2 \sim B_2 \sim C_2 \sim D_2 \sim \cdots$. □

Lemma 1.2. *$X_{H,W}$ is bounded if and only if H is bounded.*

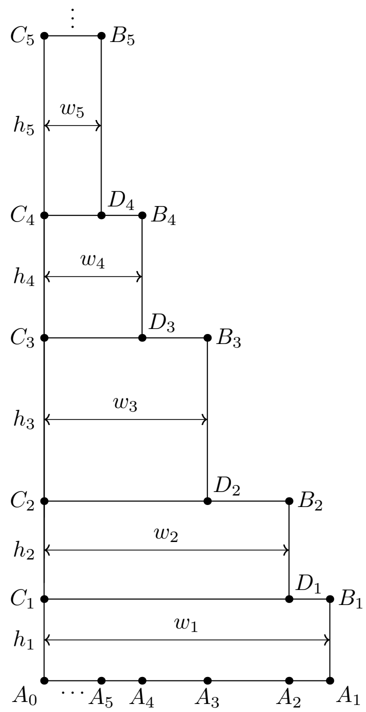

FIGURE 1. A sample surface $X_{H,W}$. In addition to the identification of vertical edges, as indicated by the arrows, each segment $A_k A_{k+1}$ is identified with $B_k D_k$ via vertical translation.

Proof. Suppose H is bounded by M_H and W by M_W. Then every point of every rectangle R_n is within $M = \sqrt{{M_H}^2 + {M_W}^2}$ of a corner. Since the vertices are identified to a single point in $\overline{X}_{H,W}$, $2M$ is an upper bound for the distance between any two points of $X_{H,W}$.

Now suppose H is not bounded. Then the centers of the rectangles R_n become arbitrarily far from the vertices, and so $X_{H,W}$ is not bounded. $\square$

Lemma 1.3. *$X_{H,W}$ is totally bounded if and only if H tends to zero.*

Proof. Let s denote the unique point in $\overline{X}_{H,W} \setminus X_{H,W}$.

Suppose H tends to zero. Then, because W also tends to zero, for every $\varepsilon > 0$ there exists N such that $h_n < \sqrt{\varepsilon}$ and $w_n < \sqrt{\varepsilon}$ for all $n \geq N$. This implies that the ε-neighborhood B_ε of s covers all R_n for $n \geq N$. The complement of B_ε in $X_{H,W}$ is compact, being a finite union of compact pieces, and can therefore be covered by finitely many ε-balls.

Suppose H does not tend to zero. Then there exists $\varepsilon_0 > 0$ such that $h_n > 2\varepsilon_0$ for infinitely many $n \geq 1$. Fix such an ε_0 and cover $\overline{X}_{H,W}$ by the following open sets:

- the ε_0-neighborhood B_{ε_0} of s;
- for each R_n not fully covered by B_{ε_0} (of which there are infinitely many), the interior of this cylinder;
- for each edge between R_n and R_{n+1}, a neighborhood of radius $(1/3) \cdot \min\{h_n, h_{n+1}\}$.

Any finite subcover of this open cover would fail to cover infinitely many interior points of the R_ns, and so $\overline{X}_{H,W}$ cannot be compact. $\square$

With these observations about $X_{H,W}$ in mind, we proceed to our counterexamples.

Example 2 (*finite area $\not\Rightarrow$ finite analytic type*)**.** Take $X_{H,W}$ with $h_n = w_n = 1/n$.

Example 3 (*finite area $\not\Rightarrow$ bounded*)**.** Take $X_{H,W}$ with $h_n = n$ and $w_n = 1/n^3$.

Example 4 (*finite area $\not\Rightarrow$ totally bounded*)**.** Take $X_{H,W}$ with $h_n = 1$ and $w_n = 1/n^2$.

Example 5 (*bounded $\not\Rightarrow$ finite analytic type*)**.** Take Example 2 or 4.

Example 6 (*bounded $\not\Rightarrow$ finite area*)**.** Take $X_{H,W}$ with $h_n = 1$ and $w_n = 1/n$.

Example 7 (*bounded $\not\Rightarrow$ totally bounded*)**.** Take Example 4 or 6.

Example 8 (*totally bounded $\not\Rightarrow$ finite analytic type*)**.** Take Example 2.

Example 9 (*totally bounded $\not\Rightarrow$ finite area*)**.** Take $X_{H,W}$ with $h_n = w_n = 1/\sqrt{n}$.

Example 10 (*finite analytic type $\not\Rightarrow$ finite area or bounded*)**.** The Riemann surface $\mathbb{C}^*$ has finite analytic type, since it is obtained from the Riemann sphere by removing two points. However, the translation structure given by the differential $\mathrm{d}z/z$ makes $\mathbb{C}^*$ isometric to an infinite cylinder, in which case it does not have finite area, and it is not bounded.

Example 10 shows that the essential way a surface of finite analytic type can fail to have finite affine type is that one could take a *meromorphic* differential on a compact Riemann surface and remove the zeroes and poles to obtain a translation surface of finite analytic type. However, if we pair the "finite analytic type" condition with any of the others, then the rest follow. This fact is likely to be well-known, but we prove it here for completeness and to show how the analytic structure of the translation surface plays a role.

Proposition 1.2. *If X has finite analytic type and finite area, then it is totally bounded.*

Proof. The translation structure is given by an abelian differential on X. Let $\widetilde{X}$ denote the compact surface from which X is obtained as a Riemann surface. Because X has finite area, the differential can be extended to $\widetilde{X}$; each point of $\widetilde{X} \setminus X$ is either a regular point or a zero of the differential. Thus $\overline{X}$ is canonically homeomorphic to $\widetilde{X}$, so X is totally bounded. □

Proposition 1.3. *If X has finite analytic type and it is bounded, then it has finite area.*

Proof. The translation structure is given by an abelian differential on X that is meromorphic on the compact Riemann surface $\widetilde{X}$ from which it is obtained by punctures. Because X is bounded, none of the punctures is at an infinite distance from any other point of X. Therefore the differential has no poles on $\widetilde{X}$, and so it has finite area. □

Proof of second part of Theorem 1. Immediate from Propositions 1.1, 1.2, and 1.3. □

To conclude this section, we observe that other collections of conditions do not imply any of the remaining ones, except as trivially follows from what has been established.

Example 11 (*finite area + totally bounded $\not\Rightarrow$ finite analytic type*)**.** Take Example 2.

Example 12 (*finite area + bounded $\not\Rightarrow$ totally bounded*)**.** Take Example 4.

2. Discreteness of Veech groups

Recently, it has become apparent that translation surfaces of infinite analytic type allow for Veech groups of much greater complexity than occurs in the case of finite type. Specifically, it is well-known that the Veech group of a translation surface of finite affine type is always a Fuchsian (i.e., discrete) subgroup of $\mathrm{SL}_2(\mathbb{R})$ and is never co-compact. In contrast, it has been shown by direct construction that any countable subgroup of $\mathrm{SL}_2(\mathbb{R})$ (in fact, of $\mathrm{GL}_2^+(\mathbb{R})$) that avoids the set of matrices with operator norm less than 1 can occur as the Veech group of a translation surface whose topological type is that of a "Loch Ness Monster", meaning it has infinite genus and one topological end [PSV]. Other "naturally occurring" examples (e.g., the surface obtained by "unfolding" an irrational polygon [Va]) also demonstrate that one cannot in general expect the Veech group of a translation surface of infinite type to be discrete. In this section, we show that this phenomenon of non-discreteness relies essentially on the failure of a surface to be totally bounded or to have finite area; i.e., it is not enough that the surface simply have infinite analytic type.

The usual proof of discreteness in the case of finite affine type is carried out by showing that the Veech group acts on the set of holonomy vectors of saddle connections, which is a discrete subset of $\mathbb{C}$ (see, e.g., [Vo]). For surfaces not of finite affine type, this last clause no longer holds: in many examples, the holonomy vectors of saddle connections do not have their lengths bounded away from zero. We find another subset of $\mathbb{C}$ on which the Veech group acts and which, under the conditions of Theorem 2, is also discrete. Our proof holds also for surfaces of finite affine type, and bypasses considerations of whether the holonomy vectors of saddle connections form a discrete set or not.

Observe, first of all, that if X has finite area, then any element of $\mathrm{Aff}^+(X)$ must preserve this area, and so the condition that $\Gamma(X) \subset \mathrm{SL}_2(\mathbb{R})$ follows automatically. Similarly, we have the following.

Lemma 2.1. *If X is totally bounded, then $\Gamma(X) \subset \mathrm{SL}_2(\mathbb{R})$.*

Proof. We use the compactness of $\overline{X}$ to establish a kind of Poincaré recurrence, which will permit us to define a first return map. Let $\phi \in \mathrm{Aff}^+(X)$. For any open subset U of X with piecewise smooth boundary, we observe that the images $\phi^{\circ n}(U)$ cannot all be disjoint: for otherwise, we could take them together with one more open subset, formed by the union of their complement and regular neighborhoods of their boundaries, and we would have an open cover of $\overline{X}$ with no finite subcover. Therefore, by a standard argument, $\phi^{\circ N}(U) \cap U \neq \varnothing$ for some $N \geq 1$. Proceeding inductively, we obtain a first return map R_ϕ into U, defined on an open subset of U whose complement has measure zero. Choose U so that it has finite area. The area of the image is

$$\mathrm{Area}(R_\phi) = \int_U (\det DR_\phi)\, d\mathrm{Area} \leq \mathrm{Area}(U).$$

If $\operatorname{der}\phi$ had determinant greater than 1, then the Jacobian determinant in the above integral would be greater than 1 on the entire domain, and the given inequality would not hold. We conclude that any element of $\mathrm{Aff}^+(X)$ must have a derivative in $\mathrm{SL}_2(\mathbb{R})$. □

Now we proceed to the main ideas in the proof of Theorem 2. Throughout this section, we take cylinders in X to be *open* subsets of X; that is, they do not include their boundaries.

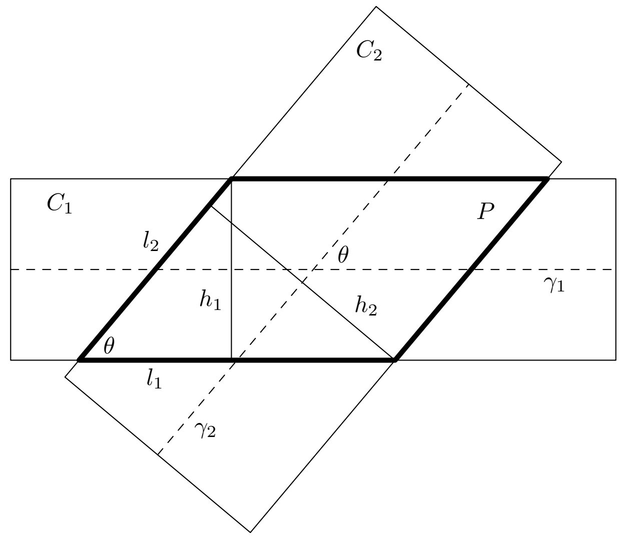

FIGURE 2. Setup for the proof of Lemma 2.2. The length of γ_1 is w_1, and the length of γ_2 is w_2.

Lemma 2.2. *Let C_1 and C_2 be two maximal cylinders in a translation surface whose respective circumferences are w_1 and w_2 and whose respective heights are h_1 and h_2, and suppose that they intersect but do not coincide. Then the angle θ between the core curves of C_1 and C_2 satisfies*

$$|\tan\theta| > \min\left\{\frac{h_1}{w_1}, \frac{h_2}{w_2}\right\}.$$

Proof. Let γ_1 and γ_2 be the core curves of C_1 and C_2, respectively. If γ_1 and γ_2 meet at right angles, then we are done. So suppose they do not meet at right angles. We note that each time γ_1 crosses one boundary component of C_2, it must cross the other boundary component before returning to the first, and likewise for γ_2 crossing the boundary of C_1. Therefore the connected components of $C_1 \cap C_2$ are Euclidean parallelograms whose sides are arcs of the boundaries of C_1 and C_2. Let P be one such parallelogram (see Figure 2). The angles of P are θ and $\pi - \theta$, so it suffices to consider the smaller of these angles. Note that h_1 and h_2 are also the two heights of P. Let l_1 and l_2 be the distances from the vertex at θ to the orthogonal projections of the adjacent vertices onto the adjacent sides of P in the directions of γ_1 and γ_2, respectively. At least one of the following inequalities holds: $l_1 < w_1$ or $l_2 < w_2$. But $\tan\theta = h_1/l_1 = h_2/l_2$, from which the desired result follows. □

Lemma 2.3. *If $v_0, v \in \mathbb{C}$ satisfy $|v_0 - v| < \varepsilon < |v_0|$, then the angle θ between v_0 and v satisfies*

$$|\tan\theta| < \frac{\varepsilon}{\sqrt{|v_0|^2 - \varepsilon^2}}.$$

Proof. Under the given conditions, the largest angle a vector v can make with v_0 is when v is tangent to the circle with radius ε centered at v_0; the assumptions imply that this angle is strictly smaller than $\pi/2$ in absolute value. The result now follows by direct calculation of the tangent of the angle in this extreme case and monotonicity of the tangent function on $(-\pi/2, \pi/2)$. □

The values h_1/w_1 and h_2/w_2 in Lemma 2.2 are, of course, the *moduli* of the cylinders. The basic idea behind the next lemma is that if two cylinders have the same area and almost the same circumference, then their moduli are not very different; we can thus play the two inequalities of Lemmata 2.2 and 2.3 against each other.

Notation. Given a translation surface X and $A > 0$, we denote by $\mathcal{C}(A)$ the set of maximal cylinders on X with area A, and by $\mathcal{V}(A) \subset \mathbb{C}$ the set of holonomy vectors of core curves of elements of $\mathcal{C}(A)$.

Lemma 2.4. *Let X be a translation surface that either is totally bounded or has finite area, and let $A > 0$. Then $\mathcal{V}(A)$ is either empty or a discrete subset of $\mathbb{C}$.*

Proof. Let $v_0 \in \mathcal{V}(A)$. If $0 < \varepsilon < |v_0|$ and $v \in \mathcal{V}(A)$ is any vector such that $|v - v_0| < \varepsilon$, then the modulus of any corresponding cylinder is bounded below by

$$f_1(\varepsilon) = \frac{A}{(|v_0| + \varepsilon)^2}.$$

On the other hand, Lemma 2.3 implies that if $|v_0 - v| < \varepsilon$ and $|v_0 - w| < \varepsilon$, then the absolute value of the tangent of the angle between v and w is bounded above by

$$f_2(\varepsilon) = \frac{2\varepsilon\sqrt{|v_0|^2 - \varepsilon^2}}{|v_0|^2 - 2\varepsilon^2}.$$

Note that, as $\varepsilon \to 0$, $f_1(\varepsilon)$ tends to $A/|v_0|^2$, while $f_2(\varepsilon)$ tends to 0. We can therefore choose $\varepsilon_0 > 0$ small enough that $f_2(\varepsilon_0) < f_1(\varepsilon_0)$. Then Lemma 2.2 implies that, for any pair of distinct elements $v, w \in \mathcal{V}(A)$ such that $|v_0 - v| < \varepsilon_0$ and $|v_0 - w| < \varepsilon_0$, the corresponding cylinders in $\mathcal{C}(A)$ must be disjoint.

If X has finite area, there can only be finitely many such cylinders, and so there can only be finitely many elements of $\mathcal{V}(A)$ within ε_0 of v_0.

If X is totally bounded, there again can be only finitely many such cylinders; otherwise we could find infinitely many disjoint balls of some fixed positive radius on X, which is impossible in a totally bounded space.

In either case, v_0 is an isolated point in $\mathcal{V}(A)$; since v_0 was arbitrary, $\mathcal{V}(A)$ is discrete. □

Lemma 2.5. *Let X be a translation surface that either is totally bounded or has finite area, and let $A > 0$. Then $\mathrm{Aff}^+(X)$ preserves $\mathcal{C}(A)$ and $\Gamma(X)$ preserves $\mathcal{V}(A)$.*

Proof. The affine image of a cylinder is a cylinder, and the maximality of a cylinder is preserved because saddle connections are sent to saddle connections by elements of $\mathrm{Aff}^+(X)$. We have already observed that an affine self-homeomorphism of X must preserve area, and so the first claim is proved. The second follows immediately. □

Lemma 2.6. *Let X be a translation surface without ideal boundary, and let $C \subset X$ be a maximal cylinder of finite area. Suppose that X is not a torus. Then the stabilizer $\mathrm{Stab}(C)$ of C in $\mathrm{Aff}^+(X)$ is a cyclic group, hence discrete.*

Proof. Because C has finite area and X is not a torus, C has an ideal boundary. Because the ideal boundary of X is empty, X does not consist only of C. Therefore each boundary component of C contains a saddle connection; call these I_1 and I_2. Any element of $\mathrm{Stab}(C)$ must also preserve the lengths of I_1 and I_2. Because the boundary of C has finite length, we may assume, up to taking a finite index subgroup, that every element in $\mathrm{Stab}(C)$ fixes I_1

and I_2. But this implies that every element of $\mathrm{Stab}(C)$ fixes the entire boundary of C and thus is a power of a full Dehn twist in C. Ergo $\mathrm{Stab}(C)$ is isomorphic to a subgroup of $\mathbb{Z}$, from which the result follows. $\square$

Proof of first part of Theorem 2. The result is already known if X is a torus, because then the Veech group is (conjugate to) $\mathrm{SL}_2(\mathbb{Z})$, so assume this is not the case.

Let γ be a periodic trajectory on X. Then the image of γ is contained in some maximal cylinder. This cylinder must have some finite area A: this is immediate if the area of X is finite, and if X is totally bounded, it follows because the points of the cylinder must be a bounded distance apart. Thus the set $\mathcal{V}(A)$ is non-empty. By Lemma 2.4, it is therefore discrete. By Lemma 2.5, $\Gamma(X)$ acts on $\mathcal{V}(A)$. So it suffices to show that the stabilizer inside $\Gamma(X)$ of a point $v \in \mathcal{V}(A)$ is discrete in $\mathrm{SL}_2(\mathbb{R})$. To see this, we observe that, up to taking a finite index subgroup, the stabilizer of v in $\Gamma(X)$ may be identified with the stabilizer inside $\mathrm{Aff}^+(X)$ of some cylinder in $\mathcal{C}(A)$. Lemma 2.6 now implies the desired result. $\square$

To prove the second part of Theorem 2, we again turn to examples.

Example 13. Let L be an "irrational" rhombus, meaning its angles are not rational multiples of π. The surface X_L obtained by unfolding L is such that $\overline{X}_L \setminus X_L$ consists of four points, arising from the vertices of L. X_L is therefore bounded. Its Veech group, however, is an indiscrete subgroup of $\mathrm{SO}(2)$, generated by rotations through the angles of L.

Example 14. The infinite cylinder of Example 10 has finite analytic type. This surface has one homotopy class of periodic trajectories; these are the images of vertical lines in the plane under the universal covering map $\zeta \mapsto e^\zeta$ from $\mathbb{C}$ to $\mathbb{C}^*$, which is made into a translation covering by taking the differential $\mathrm{d}\zeta$ on the domain. The parabolic map $(x, y) \mapsto (x, y+tx)$ of $\mathbb{R}^2 \cong \mathbb{C}$ is affine with respect to $\mathrm{d}\zeta$ for any $t \in \mathbb{R}$, and it descends to $\mathbb{C}^*$ as an affine map with respect to $\mathrm{d}z/z$, acting as a Dehn twist on each annulus $\{2\pi k \leq t \log|z| \leq 2\pi(k+1)\}$, $k \in \mathbb{Z}$. The Veech group of $(\mathbb{C}^*, \mathrm{d}z/z)$ therefore contains a copy of $\mathbb{R}$, and so it is not discrete.

Remark. It is still not known whether a surface of infinite genus that has finite area or is totally bounded can have a lattice Veech group—in particular, whether the Veech group of such a surface can be co-compact.

Acknowledgements. The author wishes to thank the Hausdorff Research Institute for Mathematics in Bonn as well as the organizers of the trimester program "Geometry and Dynamics of Teichmüller Spaces", where much of this work was carried out. Thanks also to Pascal Hubert, Gabriela Schmithüsen, and Ferrán Valdez for helpful conversations and feedback, and to the referee for useful suggestions.

References

[Bo] Joshua P. Bowman. The complete family of Arnoux–Yoccoz surfaces. To appear in *Geometriae Dedicata*.

[CGL] Reza Chamanara, Frederick P. Gardiner, and Nikola Lakic. A hyperelliptic realization of the horseshoe and baker maps. *Erg. Theory Dynam. Systems* **26** (2006), 1749–1768.

[De] Laura DeMarco. The conformal geometry of billiards. *Bull. A.M.S.* **48** (2011), 33–52.

[EG] Clifford J. Earle and Frederick P. Gardiner. Teichmüller disks and Veech's $\mathcal{F}$-structures. *Contemp. Math.* **201** (1997), 165–189.

[GJ] Eugene Gutkin and Chris Judge. Affine mappings of translation surfaces: Geometry and arithmetic. *Duke Math. J.* **103** (2000), 191–213.

[HHW] W. Patrick Hooper, Pascal Hubert, and Barak Weiss. Dynamics on the infinite staircase. To appear in *Discrete and Continuous Dynamical Systems – Series A*.

[HM] John H. Hubbard and Howard Masur. Quadratic differentials and foliations. *Acta Math.* **142** (1979), 221–274.

[HLT] Pascal Hubert, Samuel Lelièvre, and Serge Troubetzkoy. The Ehrenfest wind-tree model: periodic directions, recurrence, diffusion. *J. reine angew. Math.* **656** (2011) 223–244.

[HS] Pascal Hubert and Thomas Schmidt. Veech groups and polygonal coverings. *J. Geom. and Phys.* **35** (2000), 75–91.

[LR] Christopher J. Leininger and Alan W. Reid. A combination theorem for Veech subgroups of the mapping class group. *Geom. Funct. Analysis* **16** (2006), 403–436.

[Mc] Curtis T. McMullen. Billiards and Teichmüller curves on Hilbert modular surfaces. *J. Amer. Math. Soc.* **16** (2003), 857–885.

[Mö] Martin Möller. Variations of Hodge structures of a Teichmüller curve. *J. Amer. Math. Soc.* **19** (2006), 327–344.

[PSV] Piotr Przytycki, Gabriela Schmithüsen, and Ferrán Valdez. Veech groups of Loch Ness monsters. *Ann. Inst. Fourier* **61** (2011), 673–687.

[Th] William P. Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bull. A.M.S.* **19** (1988), 417–431.

[Va] Ferrán Valdez. Veech groups, irrational billiards and stable abelian differentials. *Discrete and Continuous Dynamical Systems – Series A* **32** (2012), 1055–1063.

[Ve] William A. Veech. Teichmüller curves in moduli space, Eisenstein series, and an application to triangular billiards. *Inv. Math.* **97** (1989), 553–583.

[Vo] Yaroslav B. Vorobets. Planar structures and billiards in rational polygons: the Veech alternative. *Russ. Math. Surv.* **51** (1996), 779–817.

IMS, Department of Mathematics, Stony Brook University, Stony Brook, NY 11794
E-mail address: joshua.bowman@gmail.com